Grids ‘98 Workshop
Research Challenges and Directions Panel[1]

W. E. Johnston

Lawrence Berkeley National Laboratory

1. This work is supported by the Director, Office of Energy Research, Office of Computation and Technology Research, Mathematical, Information, and Computational Sciences Division, of the U. S. Department of Energy under Contract No. DE-AC03-76SF00098 with the University of California.

wejohnston@lbl.gov

- **Automatic annotation and cataloguing of data is critical for on-line instruments.**
  - **short term - need a few good examples**
  - **medium term - integrate work in electronic notebooks**
- **Very high performance, distributed, data-intensive systems must use a combination of network-level, platform-level, and component-level parallelism and pipelining.**
  - **short term - need some good examples, reliable pthreads**
- **Comprehensive and "precision" monitoring is essential to widely distributed, high speed system for algorithm analysis, problem diagnosis, and dynamic configuration adaptation.**
  - **short, medium term - deploy in testbeds and applications to validate**

- **Active, autonomous resource management is essential for reliability. Basic management functions - automated server restart, data migration, congestion avoidance, etc. - can be accomplished with the same type of agent infrastructure as monitoring.**
  - **medium term - more work with agent architectures needed**
- **There will be many "middleware" systems - the "real" issue is to identify the underlying infrastructure services that are essential to all of them.**
  - **medium, long term**
- ***Routine access to multi-site, wide area testbeds is essential since almost every characteristic of such environments is different from a LAN environment.***

- **High-speed, transparent access to mass storage systems (especially HPSS) at multiple sites is a basic Grid service, because such systems are essential for managing the massive amounts of data from instruments and simulation codes.**
  - **medium, long term - a very slow-to-change community**
- **Distributed management of multi-site, multi-stakeholder resource use-conditions to provide fine-grained, "transparent" strong access control**
  - **short, medium term**

- **All resources must support "contract" scheduling so that all components are simultaneously available to build "just-in-time" systems - MSS tape marshaling and staging, network bandwidth between all resources, cache space, etc. (co-scheduling in-the-large)**
  - **short, medium, long term - depending on resource**
- **Resource scheduling "contracts" must be based on automatically analyzed and enforced site-specific policy.**
- **All resources must provide "fair sharing" mechanisms so that contract schedules can't lock everyone else out (unless that is the policy).**
  - **short, medium, long term - depending on resource**

- **The infrastructure must be protected: The grid will be an important component of the scientific environment, and will be subject to the same denial of service, theft of service, compromise of integrity and/or privacy of information, and general harassment as our general computing environment. While deployment of a grid cannot wait on the resolution of all of these issues, some of them must be addressed in order to protect the basic grid infrastructure.**
  - **short-medium term**
    - **how can deployment of secure DNS and selective use of IPsec address these issues without impeding performance?**
    - **TLS vs. SSL**

Technical Approach

From the application's point of view:

A layered collection of middleware services that provide to applications uniform views of distributed resource components and the mechanisms for assembling them into systems.

From the grid implementer's point of view:

These services extend both "up and down" through the various layers of the computing and communications infrastructure. I.e., they will "drill through" layers when necessary to achieve the required performance.

Technical Approach

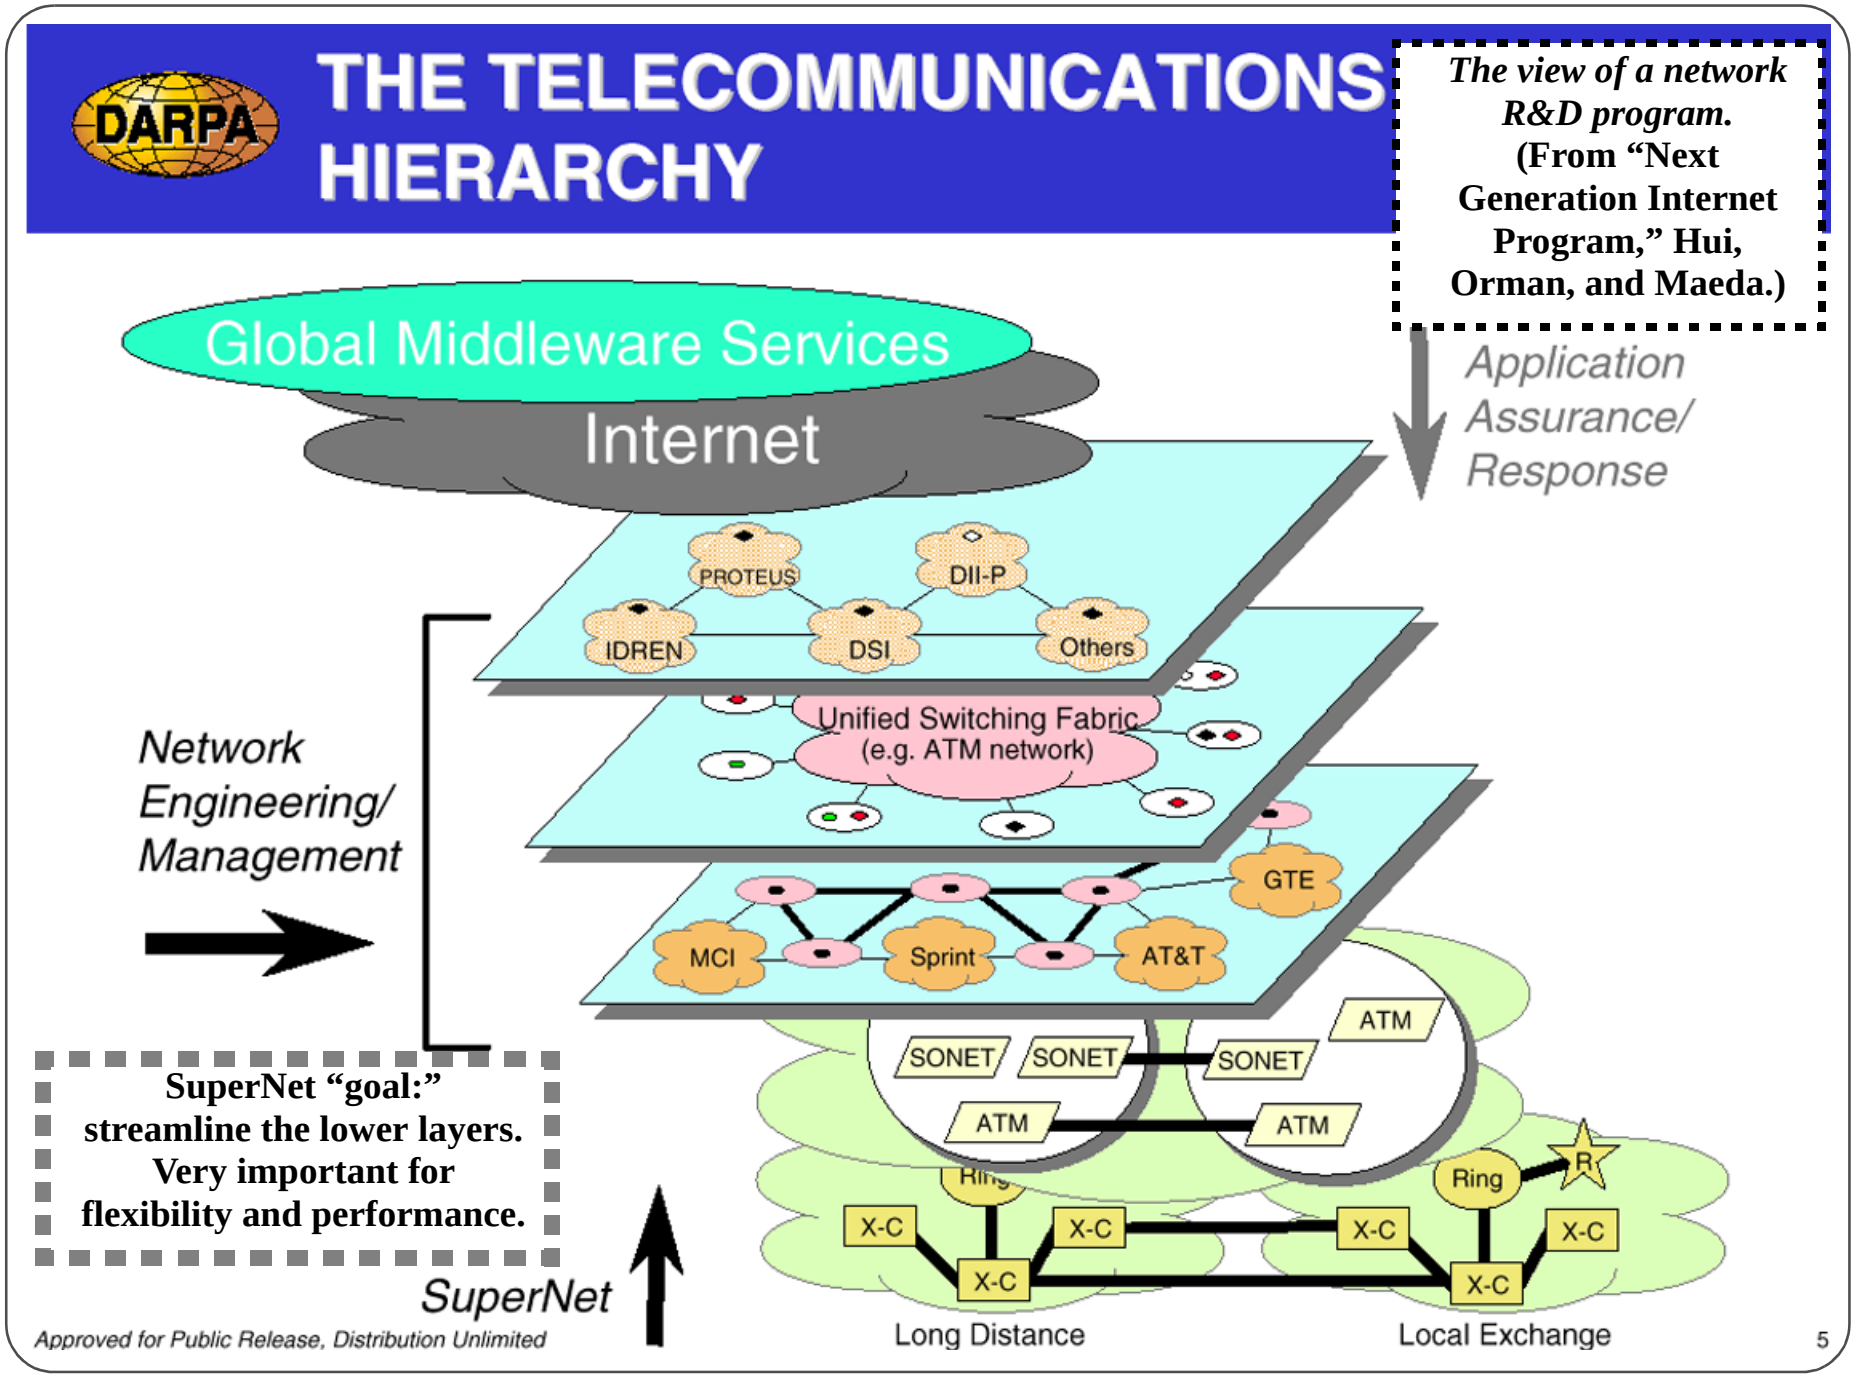

Technical Approach

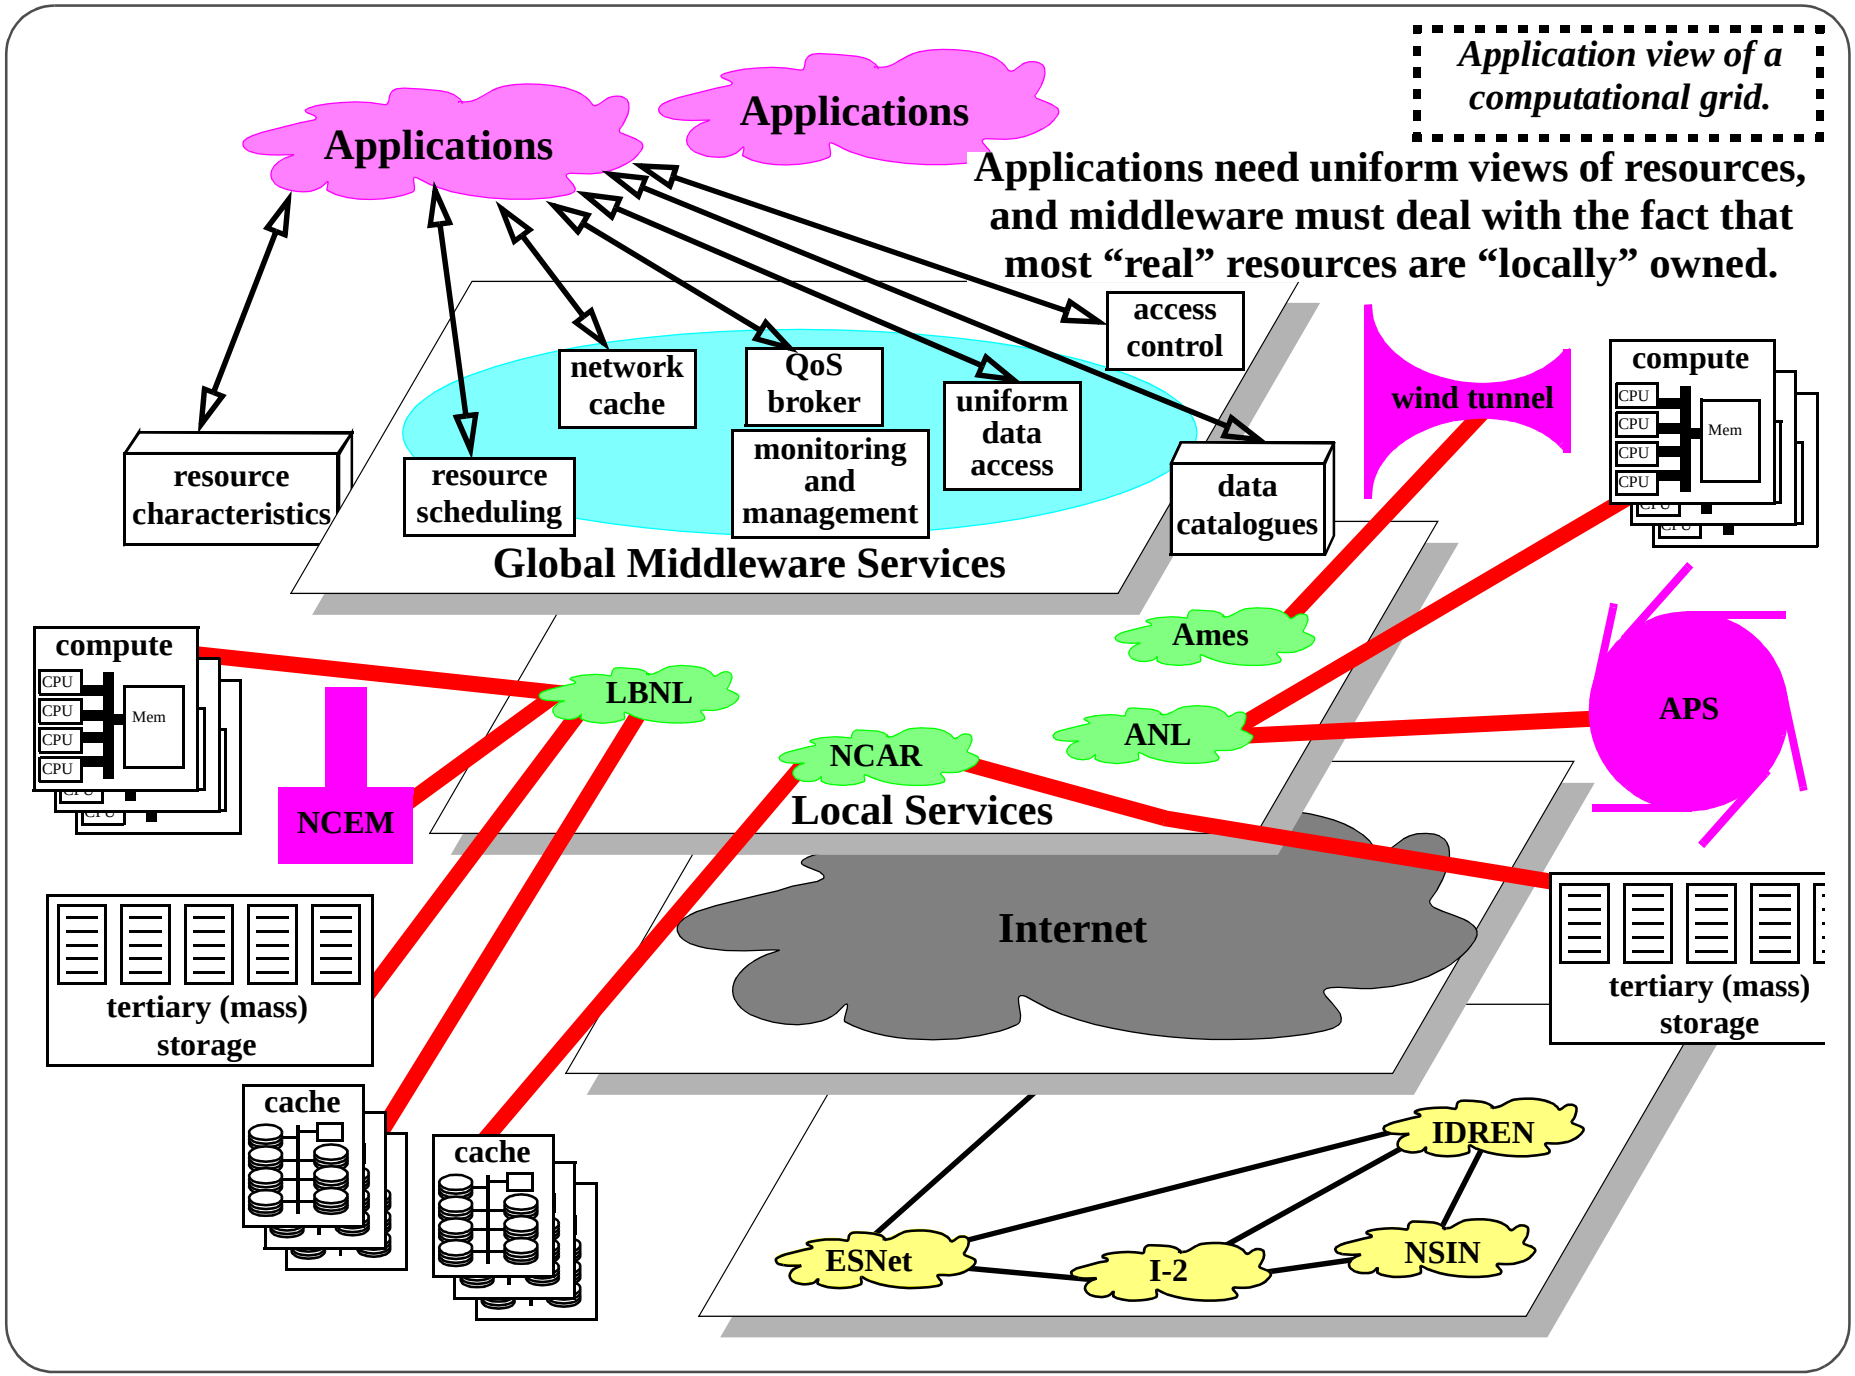

Technical Approach

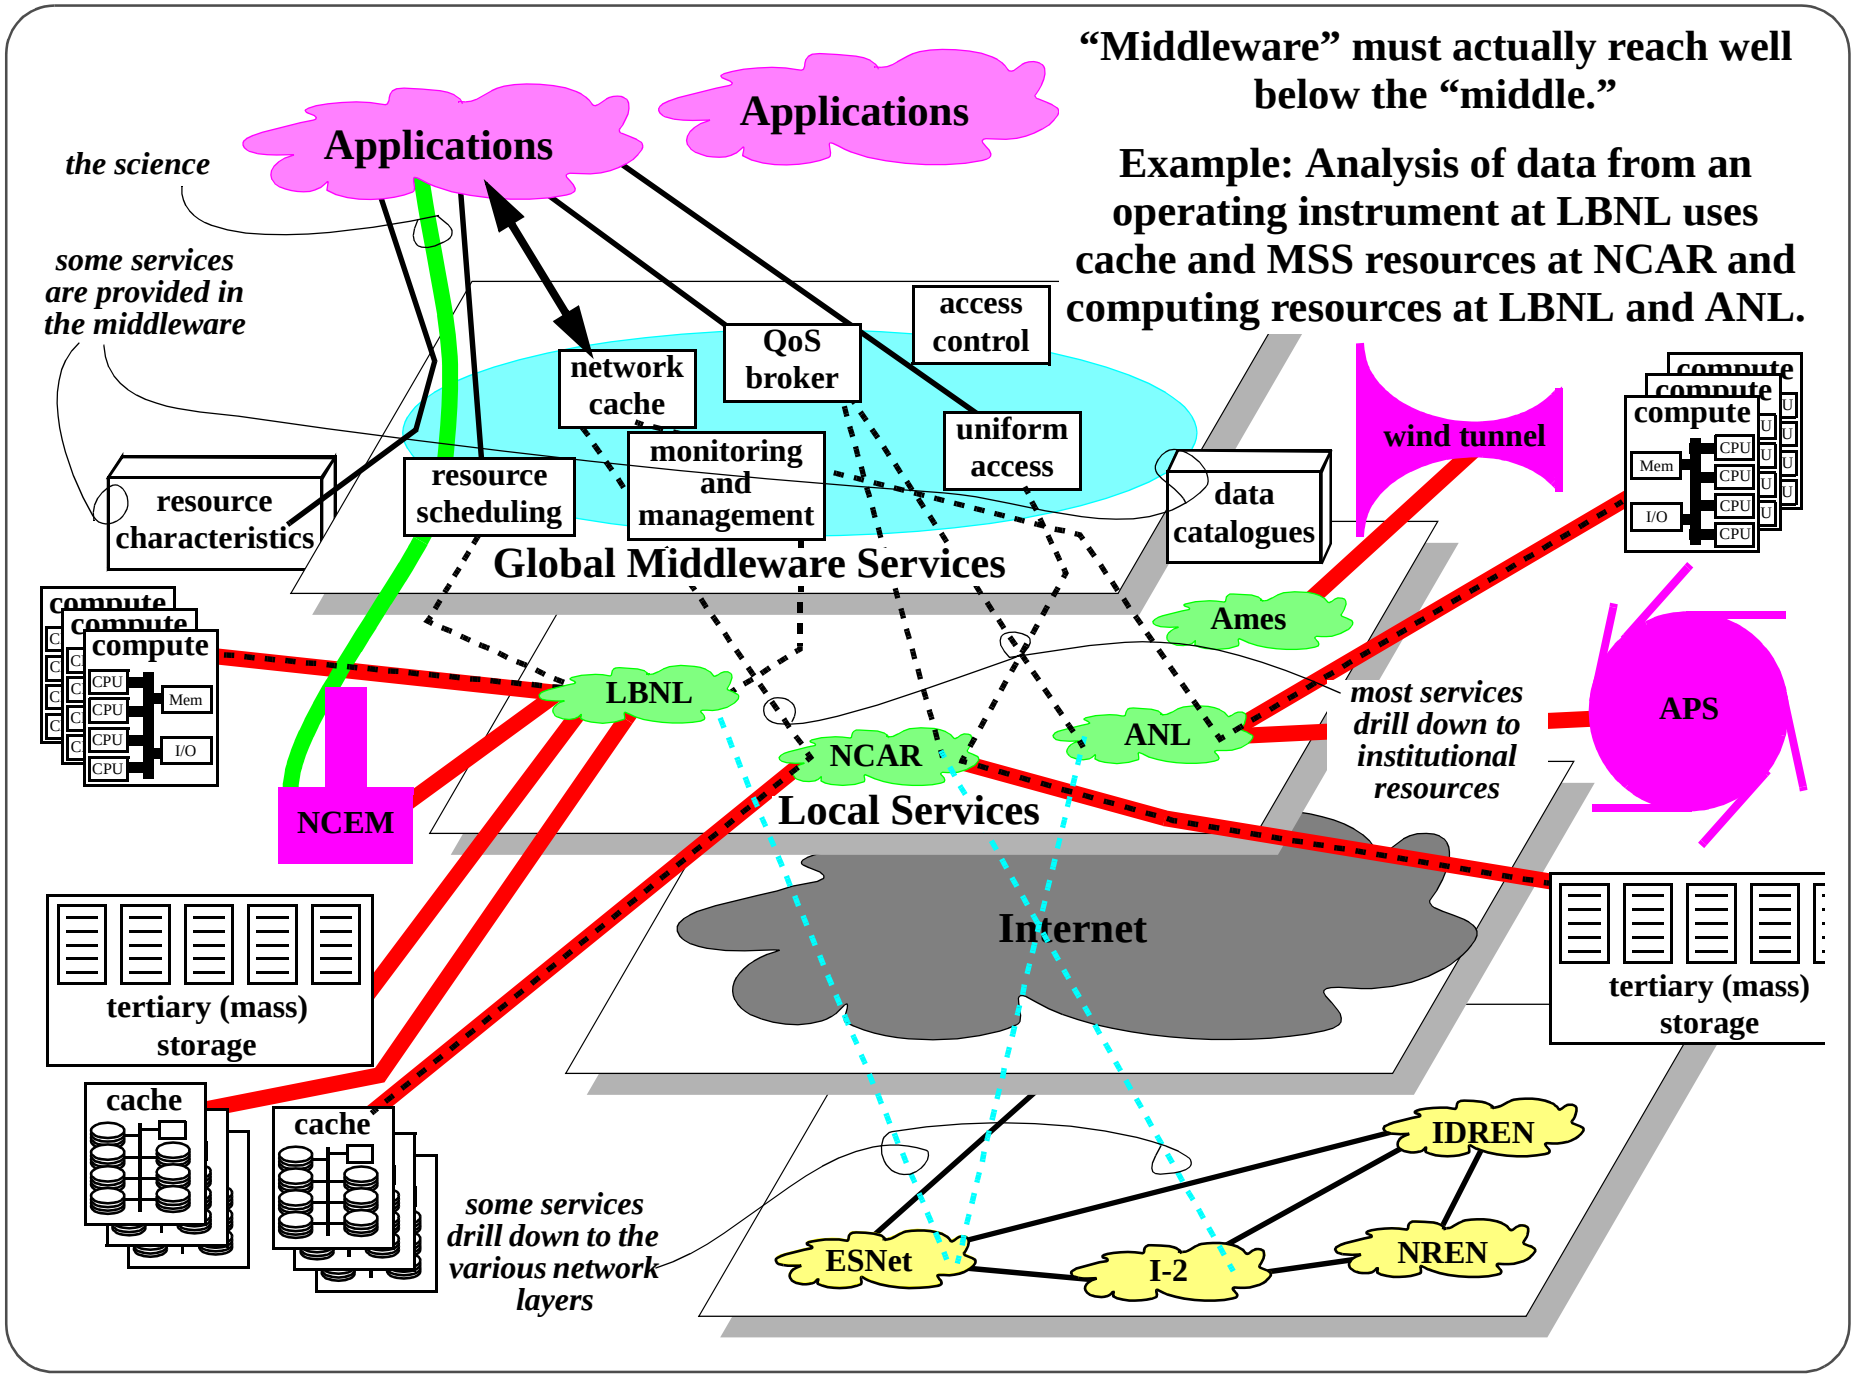

Technical Approach

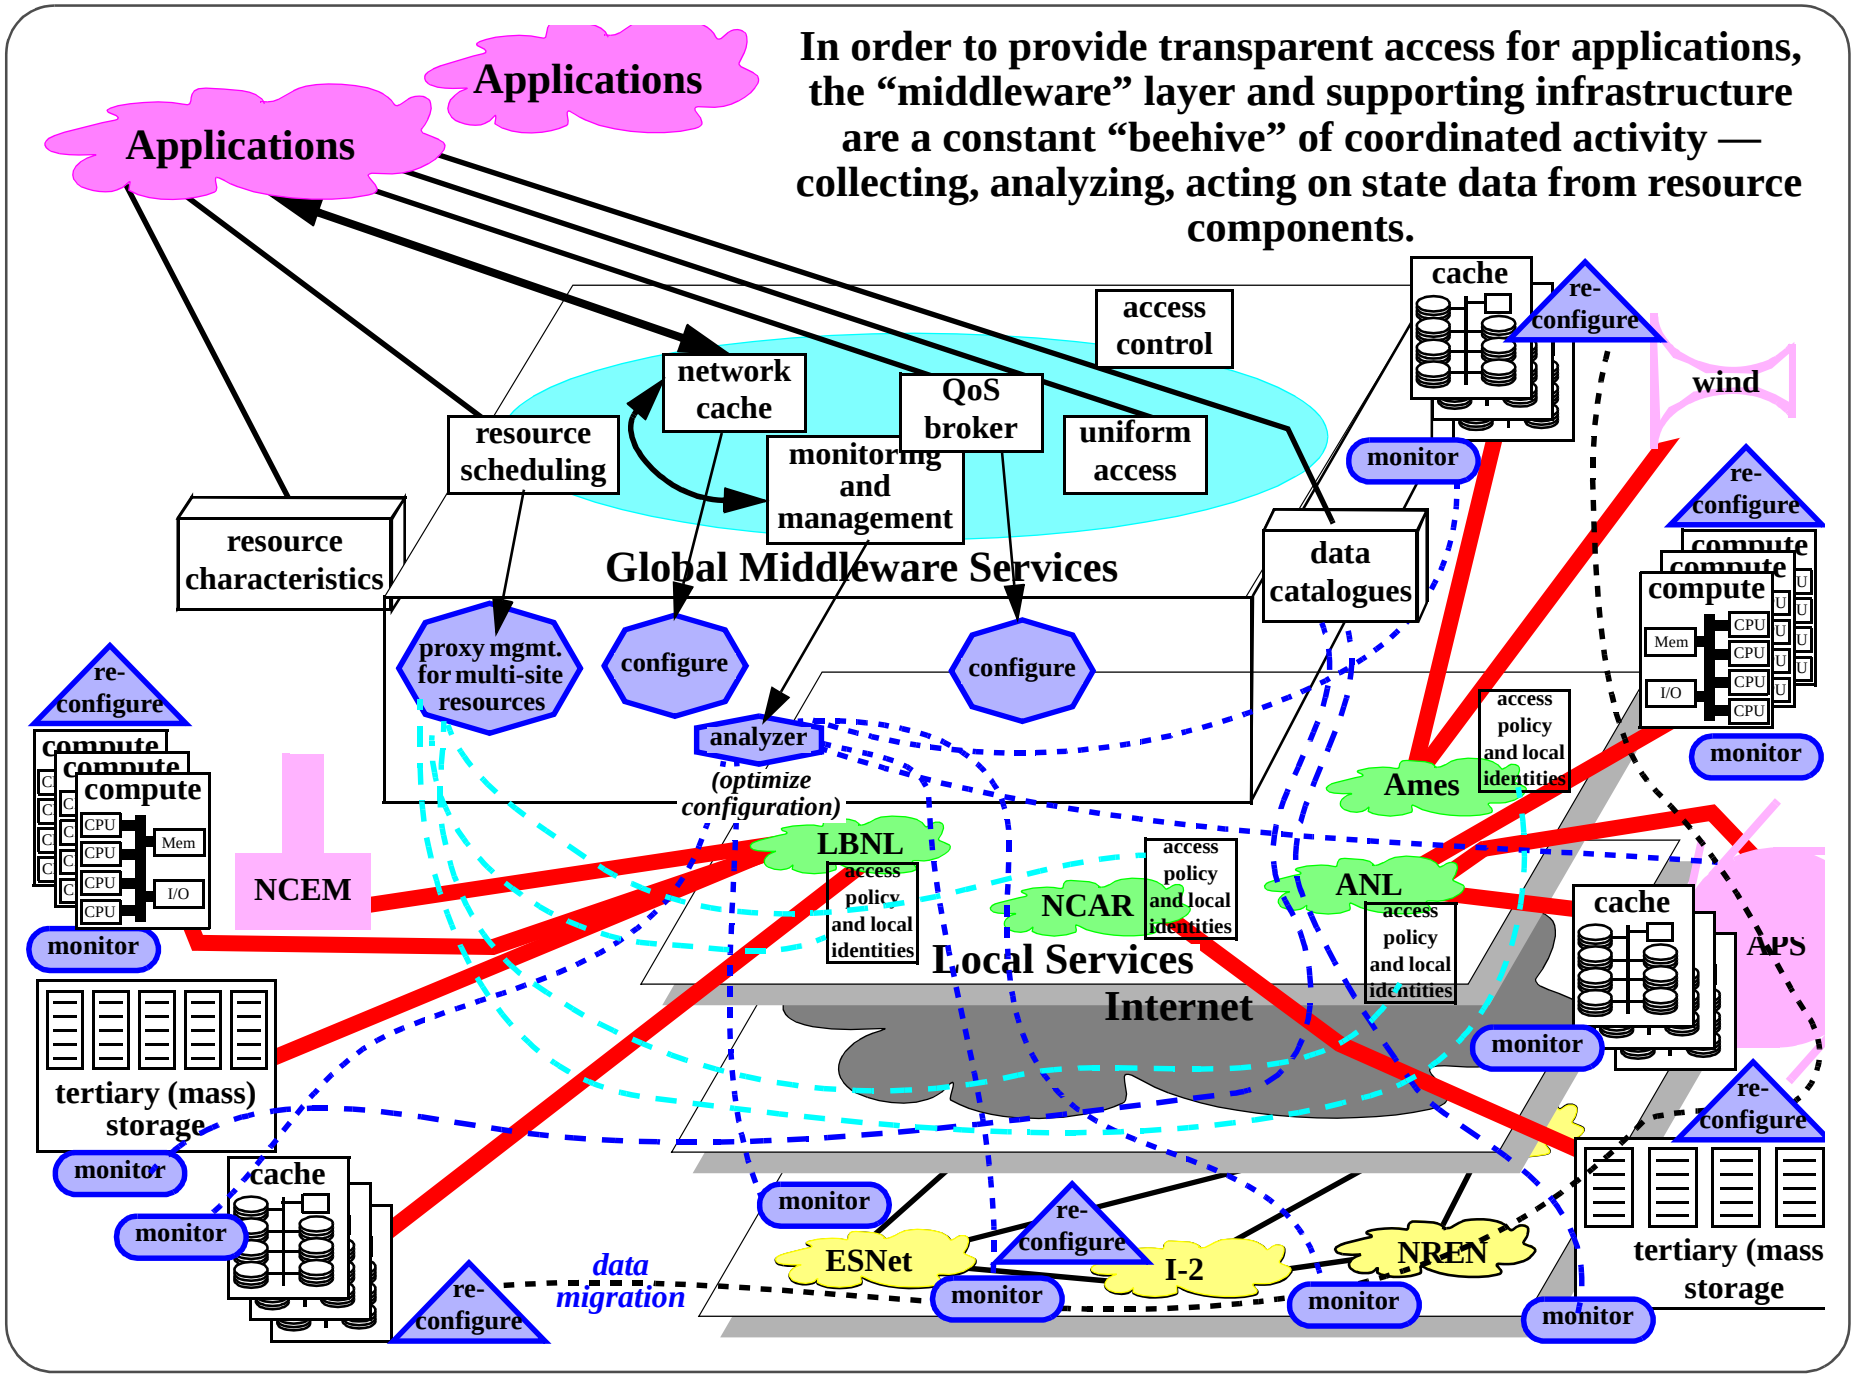

Technical Approach

The Middleware Issue

What problem is being solved (either implicitly or explicitly) by a particular middleware model?

- **Multi-disciplinary data analysis based on catalogued datasets (e.g. MDAS)**
  - **large number of unrelated users**
  - **substantial effort in metadata, etc.**
- **Single discipline analysis of very large (and growing) well structured data sets (e.g. STAF / HENP)**
  - **e.g., data from a single class of instruments**
  - **homogeneous user community (e.g. high energy physics)**
- **Automated acquisition and cataloguing of metadata-rich data sets (e.g. WALDO)**
  - **imaging systems**
  - **mixed, but limited user community**

Technical Approach

- **"Second tier" communities that use storage and computing resources wherever they can get them (e.g. Globus)**
  - **everything is scattered**
  - **users = much of the scientific community that does modeling**

Systems such as MDAS, Globus, WALDO, and STAF are all "tuned" for a particular model of resource usage (i.e. a particular community).

While there are many common elements in the middleware models, each system emphasizes different elements. Some elements may be complicated and of great interest to one community, and of little or no interest to another, who will therefore be unwilling to bear the cost or complexity. Similarly, certain models may lack elements of critical interest to a community that then sees this model as useless.

Technical Approach

There will not be a single model for defining, structuring, and implementing middleware.

However, there probably will be a set of infrastructure services can be used to support the construction of various middleware systems.

The components of the Clipper project are probably most appropriately called infrastructure services for building middleware.

Technical Approach